

The University of Chicago  
Center for Integrating Statistical and Environmental Science  
[www.stat.uchicago.edu/~cises](http://www.stat.uchicago.edu/~cises)



Chicago, Illinois USA

## **TECHNICAL REPORT NO. 37**

# **SEQUENTIAL STATE AND VARIANCE ESTIMATION WITHIN THE ENSEMBLE KALMAN FILTER**

Jonathan Stroud and Thomas Bengtsson

August 2006



Although the research described in this article has been funded wholly or in part by the United States Environmental Protection Agency through STAR Cooperative Agreement #R-82940201 to The University of Chicago, it has not been subjected to the Agency's required peer and policy review and therefore does not necessarily reflect the views of the Agency, and no official endorsement should be inferred.

# Sequential State and Variance Estimation within the Ensemble Kalman Filter

Jonathan R. Stroud\*, Thomas Bengtsson\*\*

\*Department of Statistics, The Wharton School, University of Pennsylvania

\*\*Statistics and Data Mining Group, Bell Laboratories

\*Corresponding author: Jonathan R. Stroud, Department of Statistics, The Wharton School, University of Pennsylvania, 400 Jon M. Huntsman Hall, 3730 Walnut Street, Philadelphia, PA 19104-6340. Email: [stroud@wharton.upenn.edu](mailto:stroud@wharton.upenn.edu).

## **Abstract**

Kalman filter methods for real-time assimilation of observations and dynamical systems typically assume knowledge of the system parameters. However, relatively little work has been done on extending state estimation procedures to include parameter estimation. Here, in the context of the ensemble Kalman filter, a Monte Carlo based algorithm is proposed for sequential estimation of the states and an unknown scalar observation variance. A Bayesian approach is adopted which yields analytical updating of the parameter distribution. Our proposed assimilation algorithm extends standard ensemble methods, including serial and square-root assimilation schemes. The method is illustrated on the Lorenz 40-variable system, and is shown to be robust to system nonlinearities, sparse observation networks, and the choice of the initial prior distribution.

# 1 Introduction

Motivated by the work of Evensen (1994) and Burgers et al. (1998), various ensemble Kalman filter (enKf) approaches have been proposed for Monte Carlo based data assimilation in high-dimensional atmospheric systems. Examples include serial assimilation algorithms (Houtekamer and Mitchell, 2001; Anderson, 2001), covariance stabilization techniques (Hamill et al., 2001; Anderson, 2006), and square-root filters (Bishop et al., 2001; Tippett et al., 2003). These enKf methods are typically based on the assumption of known or pre-specified system parameters, e.g., observation and background error variances, correlation ranges, and model bias terms, and the filter updating recursions are implemented conditional on fixed parameter settings. However, one may also wish to jointly estimate the states and the unknown parameters simultaneously. To this end, we propose a Bayesian algorithm for sequential estimation of unknown scalar variance parameters within the enKf framework.

In previous work addressing the issue of parameter estimation in atmospheric systems, Dee (1995) and Dee and da Silva (1999) provide a method for sequential and off-line estimation of parameters in the background and observation error covariance matrices. Their approach specifies a conditional Gaussian density function for the observational quantities given the atmospheric state and the unknown parameters, and estimates parameters by a recursive maximum likelihood (ML) procedure. An extension of this method is considered by Mitchell and Houtekamer (2000), who propose a ML-algorithm for estimation of unknown covariance parameters within the enKf framework. Another approach within the enKf context is considered by Anderson (2001), who estimates a constant forcing term by augmenting the state vector to include the unknown parameter. In recent work, Evensen (2005) considers sequential estimation when the state and parameters are jointly Gaussian.

In this paper, we take a fully Bayesian approach to the sequential estimation problem, where beliefs about the unknown parameters - along with beliefs about the system state - are represented through a joint probability distribution. This method provides a natural mechanism for incorporating prior knowledge about the parameters, and directly produces uncertainty measures for the parameters from the posterior distribution. In the case of estimating variance parameters, prior information can be based on knowledge of measurement instruments and the natural constraint of positivity. Furthermore, parameter uncertainty is allowed to impact prediction of the state as well as the observable quantities. We note that the methods of Anderson (2001) and Evensen (2005) incorporate parameter uncertainty in the state estimation process, but are based on linear updating rules and do not apply to the case of variance parameters. This work also delineates why such scale parameters cannot be efficiently estimated using linear updating or state augmentation.

To implement our method, an inverse-gamma prior distribution is assumed for the unknown variance parameter, and, conditional on the parameter, a Gaussian prior is specified for the state vector. This conditioning allows a straightforward extension of existing enKf methods (including the serial and square-root assimilation techniques) to estimate unknown scalar variance parameters. To our knowledge, this approach has not been considered in the data assimilation literature, but leads to a solution which is closely related to that of Dee (1995). The developments are given in the context of linear (or linearized) system dynamics and Gaussian forecast distributions, but simulations demonstrate robustness to system non-linearities and non-Gaussian forecast distributions.

The paper is outlined as follows. Section 2 gives notation and background for combined state and parameter estimation in state-space models. In Section 3 we consider linear Gaussian models with an unknown scalar variance. We state the form of the joint posterior density, and propose a new, ensemble-based algorithm to sample from this distribution. Section 4 provides detailed simulations to evaluate the performance of the algorithm as a function of ensemble size, system non-linearity, choice of prior distribution, and density of the observation network. Bias considerations and extensions are handled in Section 5. Section 6 provides conclusions. Serial and square-root versions of the algorithm are given in the Appendix.

## 2 Notation and Background

Let  $\mathbf{y}_t$  and  $\mathbf{x}_t$  be  $n$  and  $p$  dimensional vectors representing the observations and the unobserved state of a system at time  $t$ . Also, let  $\theta$  be a set of unknown parameters. The general state-space model is specified hierarchically through the observation and transition densities,

$$\begin{aligned} \text{Observation: } \mathbf{y}_t &\sim p(\mathbf{y}_t|\mathbf{x}_t, \theta), \\ \text{Transition: } \mathbf{x}_t &\sim p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta). \end{aligned}$$

Here,  $p(\mathbf{y}_t|\mathbf{x}_t, \theta)$  specifies the conditional probability density of the observations given the state and the parameters, and  $p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta)$  specifies the state transition density. In this work we also specify a joint prior density for the initial state and the parameters,  $p(\mathbf{x}_0, \theta|\mathbf{Y}_0)$ , where  $\mathbf{Y}_0$  denotes the information at time  $t = 0$ . Further, let  $\mathbf{Y}_t \equiv \{\mathbf{Y}_0, \mathbf{y}_1, \dots, \mathbf{y}_t\}$  denote the information up to time  $t$ .

For our problem we consider the scenario where both the observation and transition densities are Gaussian, and where the unknown parameter to be estimated is an observation noise variance. The goal of our analysis is to sequentially (in time) estimate  $\mathbf{x}_t$  and  $\theta$  through

the joint posterior distribution  $p(\mathbf{x}_t, \theta | \mathbf{Y}_t)$ . It is sometimes possible to jointly estimate  $\mathbf{x}_t$  and  $\theta$  by augmenting the state to include the parameters and then applying the Kalman filter to the augmented state vector. However, as mentioned previously, this estimation technique is only useful when the observations are linear (or approximately linear) in the unknown parameters, and cannot be implemented to estimate e.g. scale (variance) parameters. That is, for state augmentation to be successful, observation equations of the form  $\mathbf{y}_t \approx \mathbf{A}\mathbf{x}_t + \mathbf{B}\theta + \boldsymbol{\varepsilon}_t$  are implied, where  $\mathbf{A}$  and  $\mathbf{B}$  denote known matrices and  $\boldsymbol{\varepsilon}_t$  is a zero-mean white noise process. With this relationship between the data and the unknown parameter we have  $\text{cov}(\mathbf{y}_t, \theta) \neq 0$ , and the linear updating rule of the enKf can be effectively applied to update the first two moments of the parameter distribution.

In the enKf framework, state augmentation has been successfully applied by Anderson (2001) to estimate a constant forcing parameter  $F$  of the Lorenz (1996) model. In his numerical experiments, it can be shown that for a small time step  $\delta$  between observations that the (linear) relationship between the data and the unknown parameter is  $\text{cov}(\mathbf{y}_t, F) \approx \delta \text{var}(F)$ , where  $\text{var}(F)$  is the second moment of the forecast distribution  $p(F | \mathbf{Y}_{t-1})$ . In contrast, we consider parameters  $\gamma$  that multiply the observation error term and result in models of the form  $\mathbf{y}_t \approx \mathbf{A}\mathbf{x}_t + \mathbf{B}\theta + \gamma\boldsymbol{\varepsilon}_t$ . For these models it is straightforward to show  $\text{cov}(\mathbf{y}_t, \gamma) = 0$ , where (without loss of generality) we have assumed  $E(\boldsymbol{\varepsilon}_t) = 0$ .

As outlined above, state augmentation within the Kalman filter framework cannot be used as a method to estimate the observation variance. Instead, to estimate scale parameters, here denoted by  $\alpha$ , we decompose the joint posterior density as the product of a conditional state density  $p(\mathbf{x}_t | \alpha, \mathbf{Y}_t)$  and a marginal parameter density  $p(\alpha | \mathbf{Y}_t)$ . In our setting, the decomposition allows us to write separate, closed form updating recursions for the state and the parameter. The conditional state density is given by

$$p(\mathbf{x}_t | \alpha, \mathbf{Y}_t) \propto p(\mathbf{y}_t | \mathbf{x}_t, \alpha) \int p(\mathbf{x}_t | \mathbf{x}_{t-1}, \alpha) p(\mathbf{x}_{t-1} | \alpha, \mathbf{Y}_{t-1}) d\mathbf{x}_{t-1}. \quad (1)$$

Assuming a fixed parameter value  $\alpha = \alpha_0$ , the Kalman filter provides a solution to the integral in (1) for linear Gaussian models, where the state forecast density  $p(\mathbf{x}_t | \alpha_0, \mathbf{Y}_{t-1})$  is given by the integral. Similarly, the enKf and its variants mimic (1) using Monte Carlo methods. The marginal parameter density is given by

$$p(\alpha | \mathbf{Y}_t) \propto p(\mathbf{y}_t | \alpha, \mathbf{Y}_{t-1}) p(\alpha | \mathbf{Y}_{t-1}). \quad (2)$$

For the setting considered here, the marginal posterior  $p(\alpha | \mathbf{Y}_t)$  is available in closed form, and, as given in Section 3, our sampling algorithm uses the enKf to sample from (1) and the inverse gamma distribution to sample from (2).

We note that the above described conditioning is an identity and holds with  $\alpha$  replaced by a general parameter  $\theta$ . However, for general  $\theta$ , the involved integrals are typically difficult to compute in practice; specifically, it is the recursion in (2) that is hard to solve analytically. Although the computational requirements are substantial when addressing the general parameter case, the difference in uncertainty measures can be significant. To see how parameter uncertainty affects state and observation prediction when the parameters are treated as random variables rather than constants, consider the well-known variance decomposition extended to our context,

$$\text{var}(\mathbf{x}_t|\mathbf{Y}_t) = \text{var}[E(\mathbf{x}_t|\theta, \mathbf{Y}_t)] + E[\text{var}(\mathbf{x}_t|\theta, \mathbf{Y}_t)], \quad (3)$$

In the above, expectations are taken with respect to the posterior  $p(\theta|\mathbf{Y}_t)$  (see, e.g., Arnold, 1990, Sec. 3.4). Since  $\text{var}(\theta|\mathbf{Y}_t) > 0$ , we note that the inequality  $\text{var}(\mathbf{x}_t|\mathbf{Y}_t) > E[\text{var}(\mathbf{x}_t|\theta, \mathbf{Y}_t)]$  is strict. That is, the marginal posterior variance of the state is always greater than or equal to the expected variance conditional on a given value of  $\theta$ . A similar decomposition holds for prediction of a future observation  $\mathbf{y}_{t+1}$ .

[Figure 1 about here]

To illustrate the effects of parameter uncertainty in the estimation process, Figure 1 shows forecast densities for an observation (left) and posterior densities for a state variable (right). The plot is for a linear Gaussian model, and the simulation setup is described in Section 4.1. The dashed curves represent densities based on the ML approach, which uses the ML estimate for an unknown scale factor, and the solid curves represent posterior densities, which integrate out the parameter uncertainty. As can be seen in both panels, the uncertainty is larger for the posterior densities than the ML case, confirming the variance decomposition given by (3).

As mentioned, the computational cost required to evaluate (2) analytically is typically quite large in the general parameter case. However, here we consider a special case of an unknown scale factor where the parameter posterior  $p(\alpha|\mathbf{Y}_t)$  is available in closed form. In the next section we introduce the model of interest and derive the analytical form of the posterior distribution. Our developments proceed by mimicking the decompositions in (1) and (2), and yield filtering algorithms for simultaneous inference of  $\mathbf{x}_t$  and  $\alpha$ .

### 3 Gaussian State-Space Models

In what follows, we are interested in the class of state-space models expressed by the two sets of equations

$$\text{Observation:} \quad \mathbf{y}_t = H(\mathbf{x}_t) + \boldsymbol{\varepsilon}_t, \quad (4)$$

$$\text{Transition:} \quad \mathbf{x}_t = M(\mathbf{x}_{t-1}) + \boldsymbol{\omega}_t. \quad (5)$$

Here,  $H(\cdot)$  and  $M(\cdot)$  are known functions. The errors  $\boldsymbol{\varepsilon}_t$  and  $\boldsymbol{\omega}_t$  are taken as serially and mutually uncorrelated Gaussian white noise sequences with covariances  $\alpha \mathbf{R}$  and  $\alpha \mathbf{Q}$ . Letting  $\mathcal{N}(\mathbf{m}, \mathbf{C})$  denote the normal distribution with mean  $\mathbf{m}$  and covariance matrix  $\mathbf{C}$ , we have

$$\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \alpha \mathbf{R}), \quad \boldsymbol{\omega}_t \sim \mathcal{N}(\mathbf{0}, \alpha \mathbf{Q}).$$

For these densities,  $\alpha$  is an unknown scale factor to be estimated, and the matrices  $\mathbf{R}$  and  $\mathbf{Q}$  are assumed known. (Note that a deterministic system is obtained by setting  $\mathbf{Q} = \mathbf{0}$ .) With these assumptions, the implied observation and transition densities (see Section 2) are  $p(\mathbf{y}_t | \mathbf{x}_t, \alpha) = \mathcal{N}(H(\mathbf{x}_t), \alpha \mathbf{R})$ , and  $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \alpha) = \mathcal{N}(M(\mathbf{x}_{t-1}), \alpha \mathbf{Q})$ .

In the next subsections, we propose a sequential Bayesian method for estimation of  $(\mathbf{x}_t, \alpha)$  in the case of a linear Gaussian system, i.e., where  $H(\mathbf{x}_t) = \mathbf{H}_t \mathbf{x}_t$  and  $M(\mathbf{x}_t) = \mathbf{M}_t \mathbf{x}_t$  for known matrices  $\mathbf{H}_t$  and  $\mathbf{M}_t$ . As will be shown, under the linearity assumption and a particular choice of prior distribution, the joint posterior  $p(\mathbf{x}_t, \alpha | \mathbf{Y}_t)$  is available in closed form. (In Section 4, we examine the robustness of the algorithm to the assumption of linearity.)

#### 3.1 Bayesian Inference for $(\mathbf{x}_t, \alpha)$

The Bayesian model is initialized with a normal-inverse gamma prior distribution for  $(\mathbf{x}_0, \alpha)$ , providing a closed form solution to the integrals in (1) and (2). The normal-inverse gamma prior is specified as a product of the densities  $p(\mathbf{x}_0 | \alpha, \mathbf{Y}_0) = \mathcal{N}(\boldsymbol{\mu}_0, \alpha \mathbf{P}_0)$  and  $p(\alpha | \mathbf{Y}_0) = \mathcal{IG}(\nu_0/2, d_0/2)$ , where  $(\boldsymbol{\mu}_0, \mathbf{P}_0, \nu_0, d_0)$  are prespecified hyperparameters, and where  $\mathcal{IG}(a, b)$  denotes the inverse gamma distribution with parameters  $a$  and  $b$  (e.g., Hogg and Craig, 1978, Section 3.3).

Prior information about the states and parameters can be incorporated through choice of the hyperparameters. For example, when  $\mathbf{x}_t$  describes a dynamical system, one might choose the prior mean and covariance  $\boldsymbol{\mu}_0$  and  $\mathbf{P}_0$  based on system climatology. If prior information about  $\alpha$  is available, it can be incorporated by choosing  $d_0/(\nu_0 + 1)$  equal to the prior estimate (best guess), and  $\nu_0$  equal to the prior sample size (strength of belief).

Strong beliefs about  $\alpha$  are expressed by choosing a large value of  $\nu_0$ , while weak beliefs are represented by a small  $\nu_0$ . A flat (i.e. uninformative) prior on  $(\mathbf{x}_t, \alpha)$  can be obtained by setting  $\nu_0 = -1$ ,  $d_0 = 0$ , and  $\mathbf{P}_0 = c\mathbf{I}$  with large  $c$ .

With these assumptions, it can be shown that the conditional state forecast density is Gaussian at each time  $t$ , i.e.,

$$p(\mathbf{x}_t|\alpha, \mathbf{Y}_{t-1}) = \mathcal{N}(\boldsymbol{\mu}_t^f, \alpha\mathbf{P}_t^f). \quad (6)$$

Once new data  $\mathbf{y}_t$  is available, Bayes theorem can be used to calculate the joint posterior density, which is again normal-inverse gamma, and factorizes as follows

$$p(\mathbf{x}_t|\alpha, \mathbf{Y}_t) = \mathcal{N}(\boldsymbol{\mu}_t, \alpha\mathbf{P}_t), \quad (7)$$

$$p(\alpha|\mathbf{Y}_t) = \mathcal{IG}\left(\frac{\nu_t}{2}, \frac{d_t}{2}\right). \quad (8)$$

With  $\boldsymbol{\Sigma}_t = \mathbf{H}_t\mathbf{P}_t^f\mathbf{H}_t' + \mathbf{R}$  denoting the innovation covariance matrix,  $\mathbf{K}_t = \mathbf{P}_t^f\mathbf{H}_t'\boldsymbol{\Sigma}_t^{-1}$  the Kalman gain matrix,  $\mathbf{e}_t = \mathbf{y}_t - \mathbf{H}_t\boldsymbol{\mu}_t^f$  the innovation, and  $s_t = \mathbf{e}_t'\boldsymbol{\Sigma}_t^{-1}\mathbf{e}_t$  the innovation sum of squares, the hyperparameters of the densities in (6) - (8) can be obtained recursively through the following set of equations:

$$\boldsymbol{\mu}_t^f = \mathbf{M}_t\boldsymbol{\mu}_{t-1}, \quad \mathbf{P}_t^f = \mathbf{M}_t\mathbf{P}_{t-1}\mathbf{M}_t' + \mathbf{Q}, \quad (9)$$

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_t^f + \mathbf{K}_t\mathbf{e}_t, \quad \mathbf{P}_t = (\mathbf{I} - \mathbf{K}_t\mathbf{H}_t)\mathbf{P}_t^f, \quad (10)$$

$$\nu_t = \nu_{t-1} + n, \quad d_t = d_{t-1} + s_t \quad (11)$$

Notice that the recursions in (9) and (10) are identical to the forecast and update steps in the standard Kalman filter, and that the inverse gamma hyperparameters  $\nu_t$  and  $d_t$  are updated in (11) by adding the number of new (scalar) observations and the innovation sum of squares.

The recursions in (11) can be written in terms of  $\nu_t$  and the posterior mode  $d_t/(\nu_t + 1)$ . When expressed in this latter form, the updating recursion is identical to the smoothing parameter  $m_*$  and the ML estimator  $\hat{\alpha}_k$  defined in Dee (1995, Eq. 47). The connection between these two approaches is not surprising: with a flat prior distribution the posterior mode is equivalent to the ML-estimate.

Based on the above developments, the next section gives a sample-based Kalman filtering algorithm for simultaneous assimilation of observations. The algorithm extends the original enKf of Evensen (1994) and Burgers et al. (1998), but includes additional steps to update the inverse gamma hyperparameters and generate the ensemble of scale factors. The algorithm produces draws from the joint posterior distribution (7)–(8). Serial and square-root versions of the algorithm are given in Appendix B.

### 3.2 Simultaneous EnKF Scheme for $(\mathbf{x}_t, \alpha)$

We now delineate the simultaneous updating algorithm, where the entire  $n$ -dimensional vector  $\mathbf{y}_t$  is used to obtain the posterior distribution  $p(\mathbf{x}_t, \alpha | \mathbf{Y}_t)$  in a single update step. In what follows, we use the superscripts  $f$  and  $u$  to indicate the forecast (prior) and update (posterior) ensemble members, and the subscript “new” to indicate reassignment of a variable. For notational simplicity, all time subscripts are omitted. The initial prior hyperparameters  $(\boldsymbol{\mu}, \mathbf{P}, \nu, d)$  are assumed to be specified by the user.

The algorithm is initialized by drawing a sample from the initial prior distribution:  $\alpha^i \sim \mathcal{IG}(\nu/2, d/2)$  and  $\mathbf{x}^{ui} \sim \mathcal{N}(\boldsymbol{\mu}, \alpha^i \mathbf{P})$ , for  $i = 1, \dots, m$ . Then, for each assimilation time  $t = 1, \dots, T$ , we iterate between the forecast and update steps as follows. First, we generate the forecast ensemble by propagating each state  $\mathbf{x}^{ui}$  through the transition equation, i.e., let  $\mathbf{x}^{fi} = \mathbf{M}\mathbf{x}^{ui} + \boldsymbol{\omega}^i$ , where  $\boldsymbol{\omega}^i \sim \mathcal{N}(\mathbf{0}, \alpha^i \mathbf{Q})$ . Then, the forecast mean and covariance are computed based on the ensemble:

$$\hat{\boldsymbol{\mu}}^f = \frac{1}{m} \sum_{i=1}^m \mathbf{x}^{fi}, \quad \hat{\mathbf{P}}^f = \frac{1}{m-1} \sum_{i=1}^m \frac{(\mathbf{x}^{fi} - \hat{\boldsymbol{\mu}}^f)(\mathbf{x}^{fi} - \hat{\boldsymbol{\mu}}^f)'}{\alpha^i}.$$

Note that the sample covariance of the forecast ensemble members essentially produces  $\widehat{\alpha \mathbf{P}^f}$ , i.e., an estimate of  $\alpha \mathbf{P}^f$ . Thus, we divide each term  $(\mathbf{x}^{fi} - \hat{\boldsymbol{\mu}}^f)$  by  $\sqrt{\alpha^i}$  in order to obtain an estimate of the unscaled covariance matrix  $\mathbf{P}^f$ . Given the forecast moments  $\hat{\boldsymbol{\mu}}^f$  and  $\hat{\mathbf{P}}^f$ , define the following sample based quantities: the innovation  $\hat{\mathbf{e}} = \mathbf{y} - \mathbf{H}\hat{\boldsymbol{\mu}}^f$ , the innovation covariance matrix,  $\hat{\boldsymbol{\Sigma}} = \mathbf{H}\hat{\mathbf{P}}^f\mathbf{H}' + \mathbf{R}$ , the Kalman gain matrix,  $\hat{\mathbf{K}} = \hat{\mathbf{P}}^f\mathbf{H}'\hat{\boldsymbol{\Sigma}}^{-1}$ , and the innovation sum of squares,  $\hat{s} = \hat{\mathbf{e}}'\hat{\boldsymbol{\Sigma}}^{-1}\hat{\mathbf{e}}$ .

With the above quantities we generate a sample from the joint posterior of  $\alpha$  and  $\mathbf{x}$  as follows. First, generate the scale factor from its posterior distribution by setting

$$\alpha_{new}^i = \left( \frac{\hat{d}}{(\hat{d} + \hat{s})\alpha^i} + \frac{\hat{s}}{(\hat{d} + \hat{s})\tilde{\alpha}^i} \right)^{-1},$$

where  $\tilde{\alpha}^i$  is a random draw from  $\mathcal{IG}(n/2, \hat{s}/2)$ . Then, conditional on the newly drawn scale factor, generate a state vector using the perturbed observation enKf scheme: i.e. first draw the perturbation  $\boldsymbol{\epsilon}^i \sim \mathcal{N}(\mathbf{0}, \alpha_{new}^i \mathbf{R})$ , and then set

$$\mathbf{x}^{ui} = \mathbf{x}^{fi} + \hat{\mathbf{K}}(\mathbf{y} + \boldsymbol{\epsilon}^i - \mathbf{H}\mathbf{x}^{fi}).$$

The two steps above provide us with draws,  $(\alpha^1, \mathbf{x}^{u1}), \dots, (\alpha^m, \mathbf{x}^{um})$ , from the joint posterior distribution. To complete the analysis step, we update the inverse-gamma hyperparameters,  $\nu_{new} = \nu + n$  and  $\hat{d}_{new} = \hat{d} + \hat{s}$ .

Slutzky’s theorem can be used to show that, as  $m$  grows large, the above algorithm provides us with a sample  $(\mathbf{x}^{u^1}, \alpha^1), \dots, (\mathbf{x}^{u^m}, \alpha^m)$  from the joint posterior distribution,  $p(\mathbf{x}_t, \alpha | \mathbf{Y}_t)$ . We illustrate the convergence with simulations in Section 4.1. The key difference between this algorithm and the perturbed observation enKf is that each state  $\mathbf{x}^{u^i}$  is generated using a different random draw  $\alpha^i$  from the posterior distribution.

One practical issue which may impact algorithm performance, particularly in the early assimilation cycles, is that, for small  $\nu_0$ , the initialization step for  $t = 0$  may produce some extremely large  $\alpha$  draws, and this will effectively downweight the data in the updating algorithm for  $p(\mathbf{x}_1, \alpha | \mathbf{Y}_1)$ . Thus, it may be better to initialize the filter with a prior estimate that produces small  $\alpha$ -draws during the early assimilation cycles.

## 4 Simulations

In this section we use simulations to verify the accuracy of the ensemble algorithm under various conditions. The first goal is to show that the error due to Monte Carlo variability diminishes as  $m$  grows. That is, for fixed  $t$ , we present simulations verifying that  $\hat{p}(\alpha | \mathbf{Y}_t)$  approaches  $p(\alpha | \mathbf{Y}_t)$  as  $m$  tends to infinity. In Section 4.1, this convergence is verified in a linear Gaussian model where  $p(\alpha | \mathbf{Y}_t)$  can be calculated analytically, thus allowing us to confirm that the sample algorithm reproduces Bayes theorem. In Section 4.2, we consider a nonlinear deterministic system where we examine the algorithm’s robustness to nonlinear dynamics, choice of prior distribution, and sparsity of the observation network. We also investigate whether the algorithm can be combined with the state augmentation approach of Anderson (2001) to estimate an unknown forcing parameter.

### 4.1 Linear Gaussian Model

We first consider a linear Gaussian state-space model of the form in (3) and (4) with observation and transition operators defined by  $H(\mathbf{x}_t) = \mathbf{H}\mathbf{x}_t$  and  $M(\mathbf{x}_t) = \mathbf{M}\mathbf{x}_t$ . The chosen transition model is a first-order vector autoregression where the states are defined at  $p$  equally-spaced locations along a linear transect. The matrix  $\mathbf{M}$  is tridiagonal with  $\gamma_1$  on the main diagonal,  $\gamma_2$  on the superdiagonal, and  $\gamma_3$  on the subdiagonal. Following Xu and Wikle (2006), we set  $\gamma_1 = .3$ ,  $\gamma_2 = .6$  and  $\gamma_3 = .1$ . Observations are taken at each time point and each location, and we set  $\mathbf{H} = \mathbf{I}$ . For the error covariances, we assume  $\mathbf{R} = \mathbf{I}$  and  $\mathbf{Q} = \sigma^2 \mathbf{I}$ , and choose  $\sigma^2 = 1$  so that the unconditional variance of the states is approximately equal to 2. The unknown scale factor  $\alpha$  is set to 4.

Data are generated at  $T = 40$  time steps and  $n = 10$  locations. We assume a mildly informative prior with hyperparameters  $\nu_0 = 20$ ,  $d_0 = 20$ ,  $\boldsymbol{\mu}_0 = \mathbf{0}$ , and  $\mathbf{P}_0 = \mathbf{I}$ . Note that the prior distribution for  $\alpha$  has a mode of 0.95 and is thus centered at the wrong value. We run the filter algorithm defined by equations (9)–(11) to obtain the analytical posterior distribution at each time step. For comparison, we also run the enKf algorithm from Section 3 using ensemble sizes of  $m = 10, 25, 100$ , and 1000.

[Figure 2 about here]

Figure 2 shows the posterior densities at time  $t = 10$  for the fifth (scalar) state variable and the scale factor  $\alpha$ , both conditional on the data  $\mathbf{Y}_{10}$ . The bold curves represent the true densities obtained through the analytical filter, and the dashed curves represent ensemble-based density estimates for various choices of  $m$ . (The density curves are obtained using a default smoother in the statistical package *R*.) As seen in the left panel, the approximating densities converge to the true density as  $m$  grows large. Similarly, in the right panel, the approximating density for the scale factor  $\hat{p}(\alpha|\mathbf{Y}_{10})$  converges to the true posterior  $p(\alpha|\mathbf{Y}_{10})$ . These plots verify that the algorithm mimics Bayes’ theorem as  $m$  grows large. Further simulations demonstrate that the rate of convergence of the ensemble-based posterior mode to the true posterior mode is  $O(1/\sqrt{m})$ , in line with consistent Monte Carlo estimators.

## 4.2 Lorenz 96 Model

We now consider the 40-variable system of Lorenz (1996) which mimics advection at equally-spaced locations along a latitude circle. The differential equations defining the time evolution of the system are given by

$$\dot{x}_{t,k} = (x_{t,(k+1 \bmod p)} - x_{t,(k-2 \bmod p)})x_{t,(k-1 \bmod p)} - x_{t,k} + F,$$

for  $k = 1, \dots, p = 40$ . Here, we set the forcing parameter  $F$  equal to 8. We note that the system equations contain quadratic non-linearities which define a non-linear transition function  $M(\cdot)$ , and also that  $\mathbf{Q} = \mathbf{0}$  (*cf.* equation (5)). The interval between observations  $\delta$  controls the degree of nonlinearity in the system evolution, with  $\delta = .05$  producing an approximately linear system map,  $M(\mathbf{x}_t) \approx \mathbf{M}_t \mathbf{x}_t$ . In contrast,  $\delta = .25$  corresponds to a forward map with significant nonlinearities, which, in turn, yields distinctly non-Gaussian forecast distributions (see Fig.2, Bengtsson et al., 2003). A numerical solver is used to propagate the system over time.

At each time  $\delta t$ ,  $t = 1, 2, 3, \dots$ , we take  $n$  noisy observations  $y_{tj}$ ,  $j = 1, \dots, n$ , of the state: i.e.,

$$y_{tj} = \mathbf{h}'_j \mathbf{x}_t + \varepsilon_{tj}, \quad \varepsilon_{tj} \sim \mathcal{N}(0, \alpha).$$

In the above,  $\mathbf{x}_t = (x_{t,1}, \dots, x_{t,40})'$ , and  $\mathbf{h}'_j = (0, \dots, 0, 1, 0, \dots, 0)$  is an indicator vector matching the observation  $y_{tj}$  with  $x_{tj}$ . We take the errors  $\varepsilon_{tj}$  to be mutually and serially uncorrelated, and set the true observation variance to  $\alpha = 4$ . Since the climatological system variance approximately equals 16, this choice for the variance parameter produces fairly noisy observations, with a ratio of observation to system standard deviation of roughly 1/2.

To deal with Monte-Carlo error, two important tuning methods are used at each assimilation point to improve the accuracy of  $\hat{\mathbf{P}}^f$ . First, covariance stabilization is performed by multiplying sample covariances with the compact correlation function of Gaspari and Cohn (1999, Eqn. 4.10) and Hamill et al. (2001), and produces the so called ‘tapered’ forecast matrix. Secondly, covariance inflation is obtained by replacing  $\hat{\mathbf{P}}^f$  by  $(1 + k)\hat{\mathbf{P}}^f$ , where  $k$  is chosen to equal  $1/m$ . The need for these tuning methods in the context of NWP are well established (e.g., Anderson and Anderson, 1999; Houtekamer and Mitchell, 2001; Hamill et al., 2001). A formal study of the effects of covariance tapering and inflation in high-dimensional systems is given by Furrer and Bengtsson (2007). In these simulations, since the focus of our work is on the estimation of  $p(\mathbf{x}_t, \alpha | \mathbf{Y}_t)$ , we determine the tapering radius  $c$  and the inflation factor  $k$  by trial and error.

The first two simulation sets consider estimation of  $\mathbf{x}_t$  and  $\alpha$  using a full observation network with  $\mathbf{H}_t \equiv \mathbf{I}$  ( $n = 40$ ) and time steps of  $\delta = .05$  and  $\delta = .25$ . For these sets, the simultaneous assimilation scheme described in Section 3.2 is used to update the state and parameter distribution. The initial prior for  $\mathbf{x}_0$  is based on the system climatology (i.e the long-run mean and covariance), obtained by taking averages across a long time-integration of the system. Three priors for  $\alpha$  are considered:  $\mathcal{IG}(1.5, 6)$ , which represents a vague prior centered at the true value of 4;  $\mathcal{IG}(15, 240)$ , a tight prior centered at the incorrect value of 16; and  $\mathcal{IG}(15, 15)$ , a mildly informative prior centered at 1.

Our performance measure for the state estimate is the time-averaged root mean squared error (RMSE). With  $\hat{\boldsymbol{\mu}}_t$  denoting the ensemble-based posterior mean vector,  $\mathbf{x}_t$  the true state of the system, and  $T$  the number of assimilation cycles (observation times), the RMSE is given by  $\sqrt{\frac{1}{40T} \sum_{t=1}^T \sum_{k=1}^{40} (\hat{\boldsymbol{\mu}}_{t,k} - \mathbf{x}_{t,k})^2}$ . For estimation of  $\alpha$ , we give the posterior mean and the corresponding 95% empirical credible interval.

[Figure 3 about here]

Based on the serial assimilation scheme, Figure 3 illustrates the performance of the algorithm for estimating  $\alpha$ . In this plot, the ensemble size is  $m = 100$  and the tapering radius parameter is  $c = 10$ . For each of the previously described priors, the six panels depict the estimated posterior mode and the 95% empirical intervals for  $p(\alpha|\mathbf{Y}_t)$  over the first 500 assimilation cycles. The results for  $\delta = .05$  and  $\delta = .25$  are presented in the left column, and the results for  $\delta = .25$  are on the right.

For  $\delta = .05$ , the 95% intervals cover the true value for most of the assimilation period under priors 1 and 3. However, for the second prior, which is highly informative and centered at the wrong value, the algorithm takes about 100 assimilation cycles before the intervals cover the true value of 4. Note the close agreement between the three posterior distributions at  $t = 500$ , indicating algorithm robustness to prior specification. The three panels on the right shows the simulation results with  $\delta = .25$ . As previously noted, the purpose of increasing the time-interval between observations is to produce distinctly non-Gaussian forecast distributions. The panels show, although the theory in Section 3 is based on the assumption of normality, that the learning algorithm still works quite for longer forecast lead-times. However, in this case, the filter takes longer to stabilize. One reason for the efficiency of the algorithm for longer lead-times may be that the mean and covariance of the forecast distribution are estimated quite accurately, and that the algorithm is based on these first two moments.

[Table 1 about here]

Table 1 summarizes the performance of the algorithm for estimating  $\mathbf{x}_t$  and  $\alpha$  as a function of ensemble size, assimilation scheme and forecast lead time. Each simulation condition is based on 1000 assimilation cycles. The results shown on the left hand side of the table are for the simultaneous assimilation scheme, while those on the right are for the serial scheme. Further, the first four rows of the table pertain to  $\delta = .05$ , and the last four rows to  $\delta = .25$ .

We note a few crucial points about the results. First, for both assimilation schemes and for  $\delta = .05$  and  $\delta = .25$ , as the ensemble size grows, the estimate of the posterior mode converges to (or toward) the true parameter value. Second, the results for the serial scheme appears to be slightly worse than those of the simultaneous scheme. Third, as expected, the results for  $\delta = .25$  are generally worse than those for  $\delta = .05$ . Fourth, in small ensembles, both assimilation schemes appear to suffer from an upward bias in the estimation of  $\alpha$ . Moreover, when comparing estimates of the posterior mode and the RMSE between the simultaneous and the serial scheme for  $m = 10$  and  $\delta = .05$ , (i.e  $\hat{\alpha} = 4.25$  vs.  $\hat{\alpha} = 4.90$ , and RMSE=.770 vs. RMSE=1.10), we see that the serial scheme has likely produced a divergent

filter. Fifth, the RMSEs for  $\mathbf{x}_t$  are approximately twice as large for  $\delta = .25$  as for  $\delta = .05$ , whereas the posterior uncertainty for  $\alpha$  is largely independent of  $\delta$ . This last point is no surprise: the posterior variance of  $\alpha$  is controlled by the number of assimilated observations which is the same for both time steps.

In the final two simulations we consider joint estimation of  $\mathbf{x}_t$ ,  $\alpha$ , and the forcing variable  $F$ . Data is generated using  $\delta = .05$  and  $\delta = .25$  and we consider both a full and a sparse observation network. The results are obtained using the simultaneous assimilation scheme with  $m = 100$  and tapering radius  $c = 10$ . The priors are taken as  $\alpha \sim \mathcal{IG}(15, 15)$  and  $F \sim \mathcal{N}(8, 1)$ . To estimate  $F$ , we use the method of state-augmentation, where  $F$  is added to the state vector and treated as a time-varying quantity. In our implementation of this method, to improve filter stability, artificial noise with a variance of  $.5/\sqrt{t}$  is added to the time ‘evolution’ of  $F$ .

[Figure 4 about here]

Figure 4 shows the filtered mean and 95% credible intervals for  $\alpha$  and  $F$  over the first 500 assimilation cycles based on a full observation network with  $\mathbf{H} = \mathbf{I}$  ( $n = 40$ ) and  $\delta = .05$  and  $\delta = .25$ . As can be seen, the 95% credible intervals for the parameters contain the true values at the end of the assimilation period for both values of  $\delta$ , indicating robustness of our algorithm to the joint estimation of  $\alpha$  and  $F$ . Again, the uncertainty in  $\alpha$  is unaffected by the time step  $\delta$ , while the posterior uncertainty for the forcing parameter  $F$  decreases significantly as we move from  $\delta = .05$  to  $\delta = .25$ . The reason for the fast learning of  $F$  in the case of larger  $\delta$  is that the forecast variance of  $F$  grows as the time between assimilations increases, thus making the observations more informative about  $F$ . However, as pointed out earlier, as the time interval between assimilations grows, the system non-linearity will force this relationship to break down.

[Figure 5 about here]

We also consider simulations using a sparse observation network where observations are taken at every third location, i.e with  $n = 13$ . Figure 5 shows the filtered means and 95% intervals for  $\alpha$  and  $F$  over the first 1000 assimilation cycles. As in the previous plot, the top panels depict results for  $\delta = .05$ , while the bottom panels show results for  $\delta = .25$ . Again, both parameters are estimated accurately, with the 95% intervals containing the true values for most of the assimilation period, and the credible intervals narrowing over time. We note that the estimates of  $F$  show a slight upward bias at the end of the assimilation period, and

exhibit a slight oscillatory behavior for  $\delta = .25$ . The latter feature also appears in Figure 9 of Anderson (2001), and is likely due to the breakdown in the linear relationship between  $\mathbf{y}$  and  $F$  for larger time steps.

## 5 Bias Considerations and Extensions

This section considers bias adjustments and discusses extensions to include temporally and spatially varying scale factors.

### 5.1 Bias adjustments

Our simulations verify that the proposed sampling algorithm reproduces Bayes theorem as the ensemble size grows large. However, as demonstrated by the simulations, for small ensembles, the effects of Monte-Carlo error results in upward bias in terms of the estimation of the (true) posterior mode of  $p(\alpha|Y_t)$ , which equals  $d_t/(\nu_t+1)$ . Here, by considering a single Bayes update step, we show that the bias is positive for fixed  $m$ .

For given data  $\mathbf{y}_t$ , the bias is caused by sampling variability in the term  $\hat{s}_t = (\mathbf{y}_t - \mathbf{H}\hat{\boldsymbol{\mu}}_t)' \hat{\boldsymbol{\Sigma}}_t^{-1} (\mathbf{y}_t - \hat{\boldsymbol{\mu}}_t)$ , our ensemble based estimate of the true innovation sum of squares  $s_t = (\mathbf{y}_t - \mathbf{H}\boldsymbol{\mu}_t)' \boldsymbol{\Sigma}_t^{-1} (\mathbf{y}_t - \mathbf{H}\boldsymbol{\mu}_t)$ . Using classical results from multivariate statistics to evaluate the expected value of  $\hat{s}_t$ , we have

$$E(\hat{s}_t) = \left( \frac{m-1}{m-n-2} \right) s_t + \left( \frac{n}{m-n-2} \right).$$

This result can be derived using properties of the F-distribution. A bias-corrected innovation sum of squares can thus be obtained by defining  $C = (\frac{n}{m-n-2})$ ,  $B = (\frac{m-1}{m-n-2})$ , and letting  $\tilde{s}_t = (\hat{s}_t - C)/B$ . It can then easily be established that  $E(\tilde{s}_t) = s_t$ . Unfortunately, the bias adjustment is not valid when the ensemble size is smaller than the dimension of the data vector  $\mathbf{y}_t$ . In the scenario where  $m < n$ , using developments similar to those of Furrer and Bengtsson (2007), matrix expansions can be used to show

$$E(\hat{s}_t) \approx s_t + \frac{1}{m} \sum_{i=1}^n \frac{\lambda_i}{\lambda_i + 1},$$

where  $\{\lambda_1, \dots, \lambda_n\}$  are the eigenvalues of  $cov(\mathbf{H}\mathbf{x}_t|\mathbf{Y}_t)$ . Of course, this result assumes knowledge of the (true) eigenvalues, but can be used to approximate the leading bias term for a given eigen-structure. For instance, for a flat spectrum, the bias of  $s_t$  is  $O(n/m)$ . However, simulations show that this result is insufficient in terms of removing the upward bias in the estimation of the posterior mode of  $\alpha$ .

The fact that the above scheme is insufficient is due to  $\hat{s}_t$  having a right-skewed distribution, and implies that a bias correction scheme which minimizes the mean-squared error (MSE) of  $\hat{s}_t$  may be more effective. Minimizing the MSE produces a multiplicative bias correction of the form  $\tilde{s}_t = (1 - k')\hat{s}_t$ , and it can be shown that  $k' \approx 1/m$ . We note that this approach yields a shrinkage type estimator similar to that obtained when estimating  $\mathbf{P}^f$  through covariance tapering. To show the effectiveness of the shrinkage-type approach, we present simulations that re-produce rows 1-2 of Table 1 with  $\hat{s}_t$  replaced by  $\tilde{s}_t = (1 - 1/m)\hat{s}_t$ . The results are given in Table 2. As can be seen, for the simultaneous assimilation scheme, the approach of minimizing the MSE of the ensemble based innovation sum of squares works quite well for  $m = 10$  and  $m = 25$ . However, although the results for  $m = 25$  are encouraging, the serial assimilation still shows sign of filter divergence for  $m = 10$ .

## 5.2 Extensions to multiple scale factors

We note that the normal inverse-gamma theory presented in Appendix A does not extend to include temporally or spatially varying scale factors. That is, it is not possible to obtain a density decomposition akin to that given by (1) and (2) which also yields a closed form solution for the posterior distribution. However, simulations show that the algorithm presented in Section 3.2, which explicitly includes a step to estimate the unscaled forecast covariance  $\mathbf{P}^f$ , converges toward the true parameter values even in settings with multiple  $\alpha$ 's. We provide two simulation examples to demonstrate this.

In the first simulation we consider temporally varying  $\alpha$  where we alternate between taking observations with scale factors  $\alpha_1 = 4$  and  $\alpha_2 = 1$ . That is, with  $y_{tj} = x_{tj} + \varepsilon_{tj}$ , we set  $\varepsilon_{tj} \sim \mathcal{N}(0, \alpha_1)$  at odd time points ( $t = 1\delta, 3\delta, \dots$ ), and set  $\varepsilon_{tj} \sim \mathcal{N}(0, \alpha_2)$  at even time points. In the second simulation, we take data from an observation network with spatially varying  $\alpha$ 's. Here, for odd spatial indices ( $j = 1, 3, \dots, 39$ ) we set  $\alpha_1 = 4$ , while for even spatial indices we set  $\alpha_2 = 9$ . Both simulation sets consider a full observation network ( $j = 1, 2, \dots, n = 40$ ). Thus, in the second setting, we take  $n_1 = 20$  observations with  $\alpha_1 = 4$  and  $n_2 = 20$  observations with  $\alpha_2 = 9$  at every time point. Data is observed every  $\delta = .05$  time points, the ensemble size is  $m = 100$ , and the tapering radius is  $c = 10$ . The simultaneous assimilations scheme is adopted, and the previously described bias correction is employed. Mildly informative priors centered at the correct parameter values are used to initialize the filter. Results from both simulation sets are shown in Figure 6.

[Figure 6 about here]

The top and bottom rows of Figure 6 show the results for the case of temporally and spatially varying scale factors, respectively. Each of the four panels show estimates of the scale factor along with the corresponding 95% credible intervals for the first  $T = 4000$  assimilation cycles. As can be seen, the algorithm works well in the case of temporally varying scale factors, but produces a more erratic filter in the case of spatially varying scale factors. We speculate that the robustness of the algorithm to temporally changing scale factors is due to the fact that the non-linear forward map of the Lorenz system provides mixing of the samples between assimilation cycles. However, in our algorithm, such mixing is not present in the case of spatially varying scale factors that pertain to observations which are assimilated simultaneously in blocks at the same time point.

## 6 Conclusions

We outline a Bayesian approach to sequential state and parameter estimation within the enKf framework. In this approach, the joint posterior distribution for the states and parameters is decomposed into two parts: the conditional posterior for the states given the parameters; and the marginal posterior for the parameters. Our work considers the special case of a linear Gaussian model with unknown scale factors, and provides an algorithm to sample from the joint posterior distribution of the parameters and the states. Most data assimilation schemes require knowledge of such parameters, and our work provides an automatic way of selecting them. We also present serial and square-root versions of the algorithm.

For the linear Gaussian model, simulations show that the algorithm converges to the analytical form of the posterior as the ensemble size grows large. Further, using the Lorenz (1996) model, the method is shown to perform well under strong departures from linearity, incorrectly centered prior distributions, and sparse observation networks. We also show that the approach can be easily combined with the state augmentation idea of Anderson (2001) to estimate an unknown forcing parameter.

In current research, we are exploring approximate Bayesian methods for estimating general covariance parameters (e.g., correlation length scales). For these parameters, no analytical form is available for the parameter posterior, requiring solutions for which heavy computational resources are needed. To deal with such computational expense in real-time systems, one may proceed by approximating the marginal posterior for the parameters, and this is the topic of future research.

## Appendix A: Normal Inverse-Gamma Results

Here we derive the distributional results given in Section 3.1; see also West and Harrison (1997, Ch. 4).

At time  $t - 1$ , we assume that the joint posterior distribution for  $(\mathbf{x}_{t-1}, \alpha)$  is normal inverse-gamma where  $p(\mathbf{x}_{t-1}|\alpha, \mathbf{Y}_{t-1}) = \mathcal{N}(\boldsymbol{\mu}_{t-1}, \alpha \mathbf{P}_{t-1})$  and  $p(\alpha|\mathbf{Y}_{t-1}) = \mathcal{IG}(\nu_{t-1}/2, d_{t-1}/2)$ , where  $\mathcal{N}(\mathbf{m}, \mathbf{C})$  denotes a normal density with mean  $\mathbf{m}$  and covariance  $\mathbf{C}$ , and  $\mathcal{IG}(a, b)$  the inverse gamma density with parameters  $a$  and  $b$ . (See Figure 7 for various  $\mathcal{IG}$  density functions.) The forecast and posterior distributions at time  $t$  are derived as follows.

For the state vector  $\mathbf{x}_t$ , the conditional forecast density is

$$\begin{aligned} p(\mathbf{x}_t|\alpha, \mathbf{Y}_{t-1}) &= \int p(\mathbf{x}_t|\mathbf{x}_{t-1}, \alpha) p(\mathbf{x}_{t-1}|\alpha, \mathbf{Y}_{t-1}) d\mathbf{x}_{t-1} \\ &= \int \mathcal{N}(\mathbf{x}_t|\mathbf{M}_t\mathbf{x}_{t-1}, \alpha\mathbf{Q}) \mathcal{N}(\mathbf{x}_{t-1}|\boldsymbol{\mu}_{t-1}, \alpha\mathbf{P}_{t-1}) d\mathbf{x}_{t-1} \\ &= \mathcal{N}(\mathbf{M}_t\boldsymbol{\mu}_{t-1}, \alpha(\mathbf{M}_t\mathbf{P}_{t-1}\mathbf{M}_t' + \mathbf{Q})) \\ &\equiv \mathcal{N}(\boldsymbol{\mu}_t^f, \alpha\mathbf{P}_t^f). \end{aligned}$$

A similar convolution yields the conditional forecast density for  $\mathbf{y}_t$ , i.e.  $p(\mathbf{y}_t|\alpha, \mathbf{Y}_{t-1}) = \mathcal{N}(\mathbf{H}_t\boldsymbol{\mu}_t^f, \alpha\boldsymbol{\Sigma}_t)$ . Further, Bayes theorem conditional on  $\alpha$  produces the posterior density for  $\mathbf{x}_t$ , i.e.  $p(\mathbf{x}_t|\alpha, \mathbf{Y}_t) = \mathcal{N}(\boldsymbol{\mu}_t, \alpha\mathbf{P}_t)$ .

For the unknown scale factor  $\alpha$ , the posterior distribution is derived as follows. The prior distribution is inverse-gamma with density  $p(\alpha|\mathbf{Y}_{t-1}) = c\alpha^{-\frac{\nu_{t-1}}{2}-1} \exp\{-\frac{d_{t-1}}{2\alpha}\}$ , where  $c$  is a normalizing constant. The likelihood is proportional to  $p(\mathbf{y}_t|\alpha, \mathbf{Y}_{t-1}) \propto \alpha^{-\frac{n}{2}} \exp\{-\frac{s_t}{2\alpha}\}$ , where  $s_t = (\mathbf{y}_t - \mathbf{H}_t\boldsymbol{\mu}_t^f)' \boldsymbol{\Sigma}_t^{-1} (\mathbf{y}_t - \mathbf{H}_t\boldsymbol{\mu}_t^f)$  denotes the innovations sum of squares. The posterior distribution is then proportional to the prior times likelihood

$$p(\alpha|\mathbf{Y}_t) = \tilde{c}\alpha^{-(\frac{\nu_{t-1}+n}{2})-1} \exp\left\{-\frac{d_{t-1} + s_t}{2\alpha}\right\},$$

with normalizing constant  $\tilde{c}$ . This can be recognized as an inverse gamma density with parameters  $\nu_t/2$  and  $d_t/2$ , where  $\nu_t = \nu_{t-1} + n$  and  $d_t = d_{t-1} + s_t$ .

## Appendix B: Serial and Square Root Assimilation

Here we derive serial and square-root versions of the algorithm presented in Section 3.2. We assume a linear Gaussian model and independent observation errors with  $\mathbf{R} = \mathbf{I}$ . (Note that the observation equation can always be “whitened” by pre-multiplication by  $\mathbf{R}^{-1/2}$ .) Letting  $j$  denote the observation index and  $\mathbf{h}'_{tj}$  the  $j$ th row of  $\mathbf{H}_t$ , the observation equation is written in scalar form as follows

$$y_{tj} = \mathbf{h}'_{tj}\mathbf{x}_t + \varepsilon_{tj}, \quad \varepsilon_{tj} \sim \mathcal{N}(0, \alpha),$$

for  $j = 1, \dots, n$ . At a fixed time  $t$ , the goal of serial assimilation is to compute the joint posterior distribution  $p(\mathbf{x}_t, \alpha | \mathbf{Y}_{tj})$  for each observation  $j = 1, \dots, n$ , where  $\mathbf{Y}_{tj}$  denotes the information up to time  $t$  and observation  $j$ .  $\mathbf{Y}_{t0}$  is the information at the start of time  $t$ .

Assuming the same prior as in Section 3.1, we obtain a normal inverse-gamma posterior for all  $t$  and  $j$ ; i.e.,  $p(\mathbf{x}_t | \alpha, \mathbf{Y}_{tj}) = \mathcal{N}(\boldsymbol{\mu}_{tj}, \alpha \mathbf{P}_{tj})$  and  $p(\alpha | \mathbf{Y}_{tj}) = \mathcal{IG}(\nu_{tj}/2, d_{tj}/2)$ . The updating recursions for the hyperparameters are as follows

$$\boldsymbol{\mu}_{t0} = \mathbf{M}_t \boldsymbol{\mu}_{t-1,n} \quad \mathbf{P}_{t0} = \mathbf{M}_t \mathbf{P}_{t-1,n} \mathbf{M}'_t + \mathbf{Q} \quad (12)$$

$$\boldsymbol{\mu}_{tj} = \boldsymbol{\mu}_{t,j-1} + \mathbf{k}_{tj} e_{tj} \quad \mathbf{P}_{tj} = (\mathbf{I} - \mathbf{k}_{tj} \mathbf{h}'_{tj}) \mathbf{P}_{t,j-1} \quad (13)$$

$$\nu_{tj} = \nu_{t,j-1} + 1 \quad d_{tj} = d_{t,j-1} + s_{tj}, \quad (14)$$

and we set  $\nu_{t0} = \nu_{t-1,n}$  and  $d_{t0} = d_{t-1,n}$ . In the above,  $e_{tj} = y_{tj} - \mathbf{h}'_{tj} \boldsymbol{\mu}_{t,j-1}$  is the innovation,  $\sigma_{tj} = \mathbf{h}'_{tj} \mathbf{P}_{t,j-1} \mathbf{h}_{tj} + 1$  the innovation variance,  $\mathbf{k}_{tj} = \mathbf{P}_{t,j-1} \mathbf{h}_{tj} / \sigma_{tj}$  the Kalman gain vector, and  $s_{tj} = e_{tj}^2 / \sigma_{tj}$  is the innovation sum of squares.

The algorithms described below mimic equations (12)–(14) using ensemble methods. To simplify notation we suppress subscripts  $t$  and  $j$  except where necessary, and use the subscript “new” to denote reassignment of a variable. At time  $t - 1$  we start with a sample  $(\mathbf{x}^1, \alpha^1), \dots, (\mathbf{x}^m, \alpha^m)$  from the posterior distribution  $p(\mathbf{x}, \alpha | \mathbf{Y}_{t-1,n})$ , along with the current inverse gamma hyperparameters  $(\nu, \hat{d})$ . Next, a sample from the forecast distribution is obtained by propagating the states through the evolution equation:

$$\mathbf{x}_{new}^i = \mathbf{M}_t \mathbf{x}^i + \boldsymbol{\omega}^i, \quad \boldsymbol{\omega}^i \sim \mathcal{N}(\mathbf{0}, \alpha^i \mathbf{Q}).$$

for  $i = 1, \dots, m$ . This provides a sample from the joint forecast distribution  $p(\mathbf{x}, \alpha | \mathbf{Y}_{t0})$ . Given the observations  $\mathbf{y}_t = (y_{t1}, \dots, y_{tn})$ , the algorithms described next allow us to update this sample to obtain draws from the joint posterior distribution.

### Algorithm 1: Serial Assimilation

*Step 0:* Start with a sample  $(\mathbf{x}^1, \alpha^1), \dots, (\mathbf{x}^m, \alpha^m)$  from the forecast distribution  $p(\mathbf{x}_t, \alpha | \mathbf{Y}_{t0})$  and the current inverse gamma hyperparameters  $(\nu, \hat{d})$ .

*Step 1 (Updating):* For each observation  $j = 1, \dots, n$ :

1. Compute the sample-based mean and covariance:

$$\hat{\boldsymbol{\mu}} = \frac{1}{m} \sum_{i=1}^m \mathbf{x}^i, \quad \hat{\mathbf{P}} = \frac{1}{m-1} \sum_{i=1}^m \frac{(\mathbf{x}^i - \hat{\boldsymbol{\mu}})(\mathbf{x}^i - \hat{\boldsymbol{\mu}})'}{\alpha^i}.$$

2. Compute the innovation, innovation variance, Kalman gain vector, and innovation sum of squares:

$$\hat{e} = y_{tj} - \mathbf{h}_{tj}' \hat{\boldsymbol{\mu}}, \quad \hat{\sigma} = \mathbf{h}_{tj}' \hat{\mathbf{P}} \mathbf{h}_{tj} + 1, \quad \hat{\mathbf{k}} = \hat{\mathbf{P}} \mathbf{h}_{tj} / \hat{\sigma}, \quad \hat{s} = \hat{e}^2 / \hat{\sigma}.$$

3. Update the inverse gamma hyperparameters:

$$\nu_{new} = \nu + 1, \quad \hat{d}_{new} = \hat{d} + \hat{s}.$$

4. For  $i = 1, \dots, m$ , draw a random variate  $\tilde{\alpha}^i \sim \mathcal{IG}(1/2, \hat{s}/2)$ , and set

$$\alpha_{new}^i = \left( \frac{\hat{d}}{(\hat{d} + \hat{s})\alpha^i} + \frac{\hat{s}}{(\hat{d} + \hat{s})\tilde{\alpha}^i} \right)^{-1}.$$

5. For  $i = 1, \dots, m$ , draw a random variate  $\varepsilon^i \sim \mathcal{N}(0, \alpha_{new}^i)$ , and set

$$\mathbf{x}_{new}^i = \mathbf{x}^i + \hat{\mathbf{k}}(y_{tj} + \varepsilon^i - \mathbf{h}_{tj}' \mathbf{x}^i).$$

## Algorithm 2: Square-Root Assimilation

*Step 0:* Start with a sample  $(\mathbf{x}^1, \alpha^1), \dots, (\mathbf{x}^m, \alpha^m)$  from the forecast distribution  $p(\mathbf{x}_t, \alpha | \mathbf{Y}_{t0})$  and the current inverse gamma hyperparameters  $(\nu, \hat{d})$ .

*Step 1 (Initialization):* Compute the forecast mean,  $\hat{\boldsymbol{\mu}}$ , and construct the  $p \times m$  perturbation matrix  $\mathbf{Z} = [\mathbf{z}^1, \mathbf{z}^2, \dots, \mathbf{z}^m]$ , where

$$\mathbf{z}^i = (\mathbf{x}^i - \hat{\boldsymbol{\mu}}) / \sqrt{\alpha^i}, \quad i = 1, \dots, m.$$

*Step 2 (Updating):* For each observation  $j = 1, \dots, n$ :

1. Compute the covariance, innovation, innovation variance, Kalman gain vector, and innovation sum of squares:

$$\hat{\mathbf{P}} = \mathbf{Z}\mathbf{Z}' / (m - 1), \quad \hat{e} = y_{tj} - \mathbf{h}_{tj}' \hat{\boldsymbol{\mu}}, \quad \hat{\sigma} = \mathbf{h}_{tj}' \hat{\mathbf{P}} \mathbf{h}_{tj} + 1, \quad \hat{\mathbf{k}} = \hat{\mathbf{P}} \mathbf{h}_{tj} / \hat{\sigma}, \quad \hat{s} = \hat{e}^2 / \hat{\sigma}.$$

2. Update the inverse gamma hyperparameters:

$$\nu_{new} = \nu + 1, \quad \hat{d}_{new} = \hat{d} + \hat{s}.$$

3. Compute the  $m \times m$  adjustment matrix

$$\mathbf{A} = (\mathbf{I} - \mathbf{Z}' \mathbf{h}_{tj} \mathbf{h}_{tj}' \mathbf{Z} / \hat{\sigma})^{-1/2}.$$

4. Update the mean vector and perturbation matrix

$$\hat{\boldsymbol{\mu}}_{new} = \hat{\boldsymbol{\mu}} + \hat{\mathbf{k}} \hat{e}, \quad \mathbf{Z}_{new} = \mathbf{Z} \mathbf{A}$$

*Step 3 (Transformation):* For  $i = 1, \dots, m$ , draw  $\alpha^i \sim \mathcal{IG}(\nu/2, \hat{d}/2)$ , define  $\mathbf{z}^i$  as the  $i$ th column of  $\mathbf{Z}$ , and obtain the updated state vector by setting

$$\mathbf{x}^i = \hat{\boldsymbol{\mu}} + \sqrt{\alpha^i} \mathbf{z}^i.$$

**Acknowledgments:** This work grows out of collaboration with the Geophysical Statistics Project (GSP) at the National Center for Atmospheric Research, Boulder, CO, USA. GSP is funded by the National Science Foundation under grants DMS 9815344 and DMS 9312686.

## References

- Anderson, J. L. (2001). An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review* **129**, 2884–2903.
- Anderson, J. L. (2006). Exploring the need for localization in ensemble data assimilation using an hierarchical ensemble filter. To appear: *Physica D: Nonlinear Phenomena* .
- Anderson, J. L. and Anderson, S. L. (1999). A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review* **127**, 2741–2758.
- Arnold, S. (1990). *Mathematical Statistics*. Prentice Hall: Englewood Cliffs, NJ.
- Bengtsson, T., Snyder, C., and Nychka, D. (2003). Toward a nonlinear ensemble filter for high-dimensional systems. *Journal of Geophysical Research* **108**(D24), 8775.
- Bishop, C., Etherton, B., and Majumdar, S. (2001). Adaptive sampling with the ensemble transform Kalman Filter. Part i: Theoretical aspects. *Monthly Weather Review* **129**, 420–436.
- Burgers, G., van Leeuwen, P. J., and Evensen, G. (1998). Analysis scheme in the ensemble Kalman filter. *Monthly Weather Review* **126**, 1719–1724.
- Dee, D. (1995). On-line estimation of error covariance parameters for atmospheric data assimilation. *Monthly Weather Review* **123**, 1128–1145.
- Dee, D. and da Silva, A. (1999). Maximum likelihood estimation of forecast and observation error covariance parameters. Part I: Methodology. *Monthly Weather Review* **127**, 1822–1834.
- Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research* **99**, 10143–10162.
- Evensen, G. (2005). The combined state and parameter estimation problem. *Submitted to Computational Geosciences* .
- Furrer, R. and Bengtsson, T. (2007). Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. To appear: *Journal of Multivariate Analysis* **98**(2).

- Gaspari, G. and Cohn, S. E. (1999). Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society* **125**, 723–757.
- Hamill, T. M., Whitaker, J. S., and Snyder, C. (2001). Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Monthly Weather Review* **129**, 2776–2790.
- Hogg, R. and Craig, A. (1978). *Introduction to Mathematical Statistics, 4th Edition*. New York: Macmillan.
- Houtekamer, P. L. and Mitchell, H. L. (2001). A sequential ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review* **129**, 123–137.
- Lorenz, E. (1996). Predictability: A problem partially solved. In: *Proceedings Seminar on Predictability*, 1–18. ECMWF: Reading Berkshire, UK.
- Mitchell, H. L. and Houtekamer, P. L. (2000). An adaptive ensemble Kalman filter. *Monthly Weather Review* **128**, 416–433.
- Tippett, M. K., Anderson, J. L., Bishop, C. H., Hamill, T. M., and Whitaker, J. S. (2003). Ensemble square-root filters. *Monthly Weather Review* **131**, 1485–1490.
- West, M. and Harrison, P. (1997). *Bayesian Forecasting and Dynamic Models*. New York: Springer, 2nd edition.
- Xu, K. and Winkle, C. K. (2006). Estimation of parameterized spatio-temporal dynamic models. *Journal of Statistical Planning and Inference* .

Table 1: Filtering results from the Lorenz 96 model after 1000 assimilation cycles. Data are generated using an observation variance of  $\alpha = 4$ , and the initial prior is  $p(\alpha|\mathbf{Y}_0) = \mathcal{IG}(1.5, 6)$ . The table shows the root mean squared error (RMSE) for  $\mathbf{x}_t$ , the posterior mode  $\hat{\alpha}$ , and 95% empirical credible intervals for  $\alpha$ . Rows 1-4 are based on  $\delta = .05$ ; Rows 5-8 are based on  $\delta = .25$ . The ensemble size is denoted by  $m$  and the tapering radius is given by  $c$ .

$m$	$c$	Simultaneous			Serial		
		RMSE	$\hat{\alpha}$	95% CI	RMSE	$\hat{\alpha}$	95% CI
10	2.5	.770	4.25	(4.20, 4.30)	1.01	4.90	(4.84, 4.94)
25	5.0	.553	4.04	(3.99, 4.08)	.580	4.11	(4.05, 4.16)
100	10.0	.476	4.02	(3.98, 4.06)	.483	4.03	(3.98, 4.08)
400	20.0	.430	4.03	(3.99, 4.08)	.420	4.02	(3.97, 4.07)
10	2.5	1.42	4.80	(4.76, 4.86)	1.53	5.64	(5.60, 5.69)
25	5.0	1.21	4.26	(4.20, 4.30)	1.22	4.60	(4.54, 4.65)
100	10.0	1.05	4.00	(3.95, 4.04)	1.06	4.23	(4.18, 4.27)
400	20.0	0.98	4.00	(3.96, 4.05)	0.98	4.11	(4.06, 4.15)

Table 2: Bias correction results. Same setup as in Table 1. Here,  $\delta = .05$  and with multiplicative bias correction factor of  $k' = 1/m$ .

$m$	$c$	Simultaneous			Serial		
		RMSE	$\hat{\alpha}$	95% CI	RMSE	$\hat{\alpha}$	95% CI
10	2.5	.769	3.96	(3.91, 4.02)	1.04	4.49	(4.44, 4.56)
25	5.0	.582	3.97	(3.92, 4.01)	.582	4.00	(3.95, 4.01)

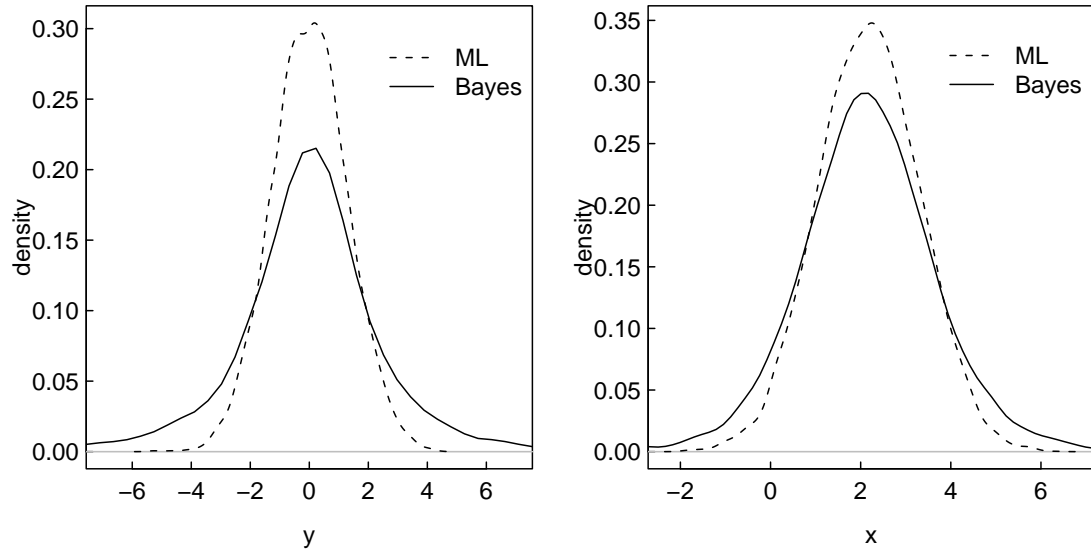


Figure 1: Forecast and posterior densities in the linear Gaussian model at time  $t = 1$ . The left panel shows the forecast densities for the fifth observation  $y_5$  under the ML approach (dashed lines) and the Bayesian approach (solid lines). The right panel shows the posterior densities for the fifth state variable  $x_5$  under the same two approaches.

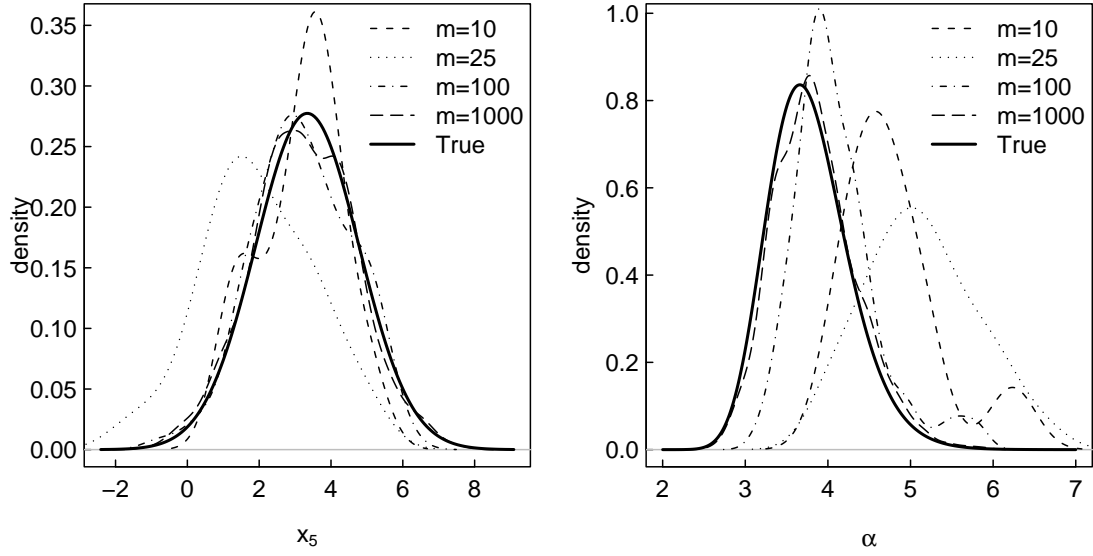


Figure 2: Posterior densities for the states and scale factor in the linear Gaussian model at time  $t = 10$ . The left panel shows the marginal posterior densities for the fifth state variable,  $x_5$ , for the ensemble-based density estimates (dashed lines) and the true density (bold line). The right panel shows the marginal posterior densities for the scale factor,  $\alpha$ , under the same approaches.

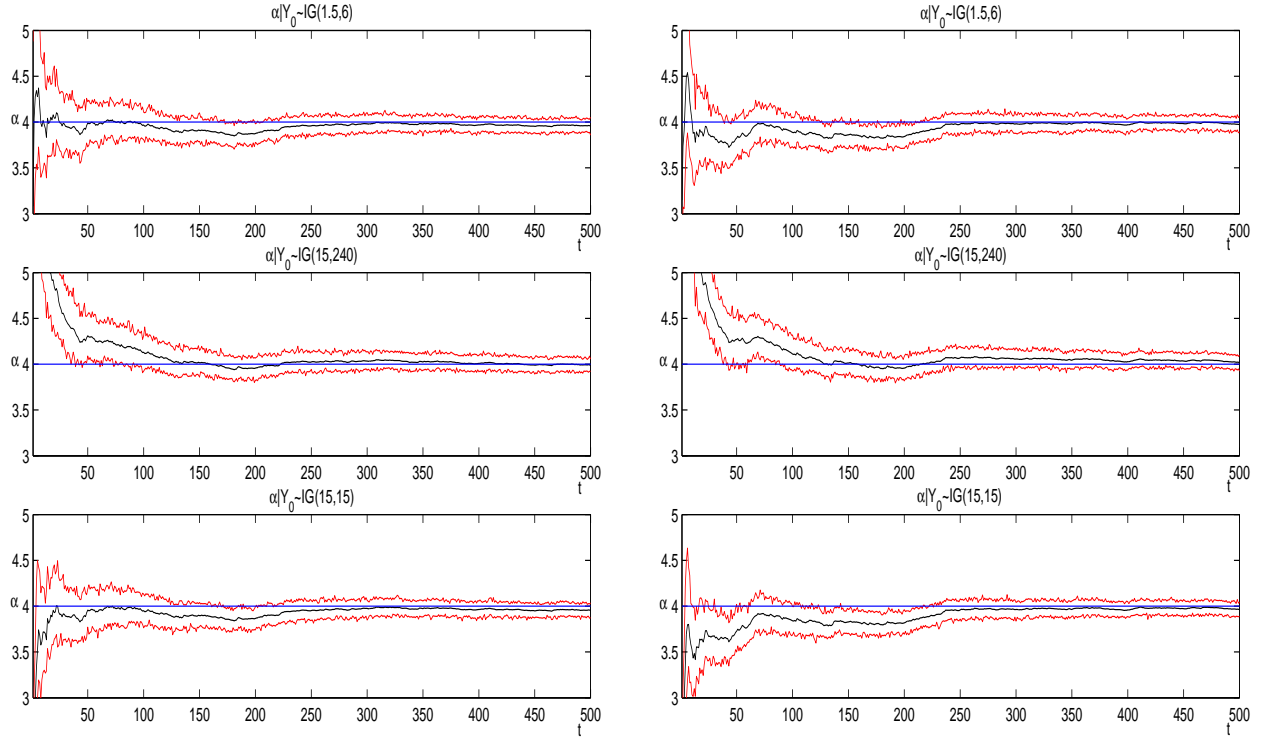


Figure 3: Posterior mode and 95% credible intervals for  $\alpha$  for the first 500 assimilation cycles under three different priors. The left panels are for  $\delta = .05$ , and the right panels are for  $\delta = .25$ . The results are based on the serial assimilation scheme, using a full observation network, an ensemble size of  $m = 100$ , and a tapering radius of  $c = 10$ .

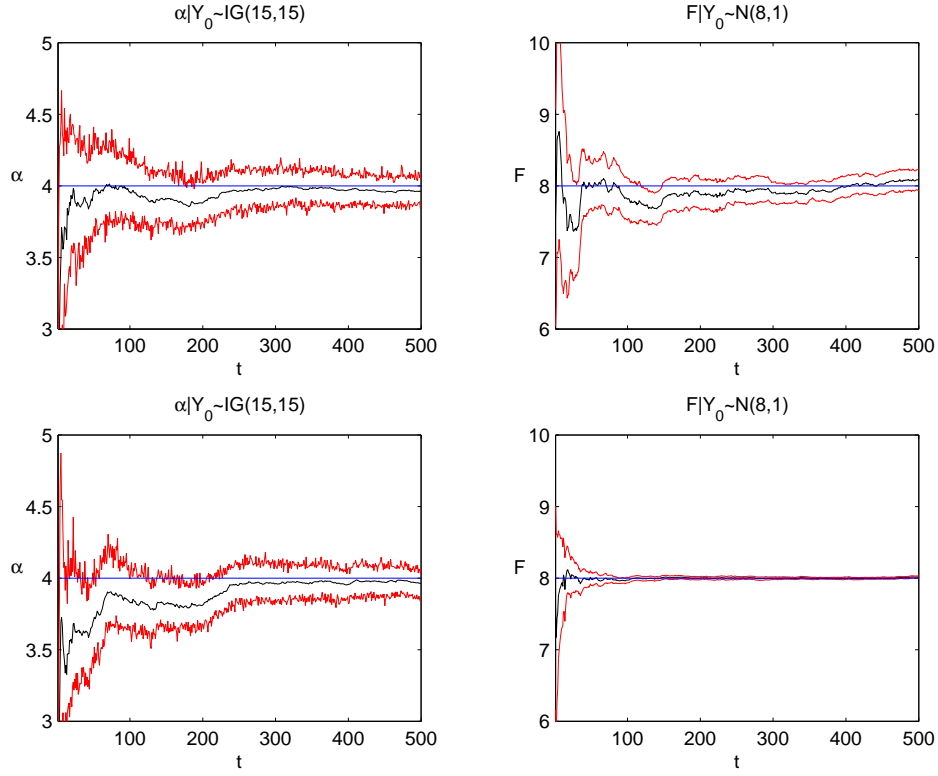


Figure 4: Posterior mode and 95% credible intervals for  $\alpha$  and  $F$  for the first 500 assimilation cycles using a full observation network. The top row is for  $\delta = .05$ , and the bottom row is for  $\delta = .25$ . The results are based on the simultaneous assimilation scheme, an ensemble size of  $m = 100$ , and a tapering radius of  $c = 10$ .

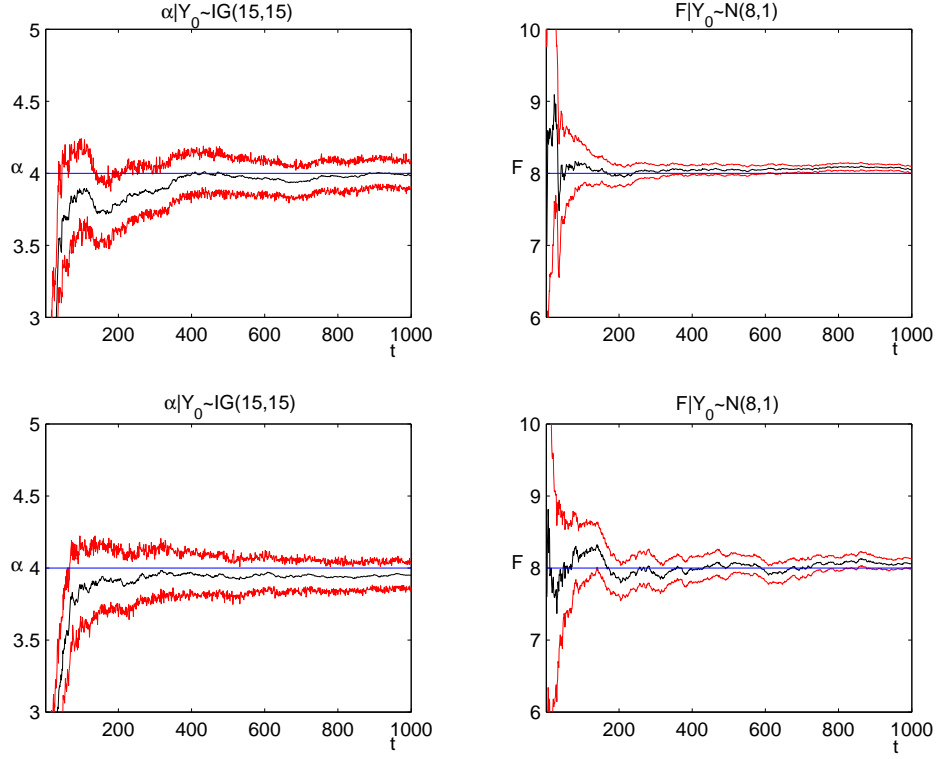


Figure 5: Posterior mean and 95% credible intervals for  $\alpha$  and  $F$  for the first 1000 assimilation cycles using a sparse observation network with observations taken at every third location,  $k = 1, 4, 7, \dots, 37$ . The top row is for  $\delta = .05$ , and the bottom row is for  $\delta = .25$ . The results are based on the simultaneous assimilation scheme, an ensemble size of  $m = 100$ , and a tapering radius of  $c = 10$ .

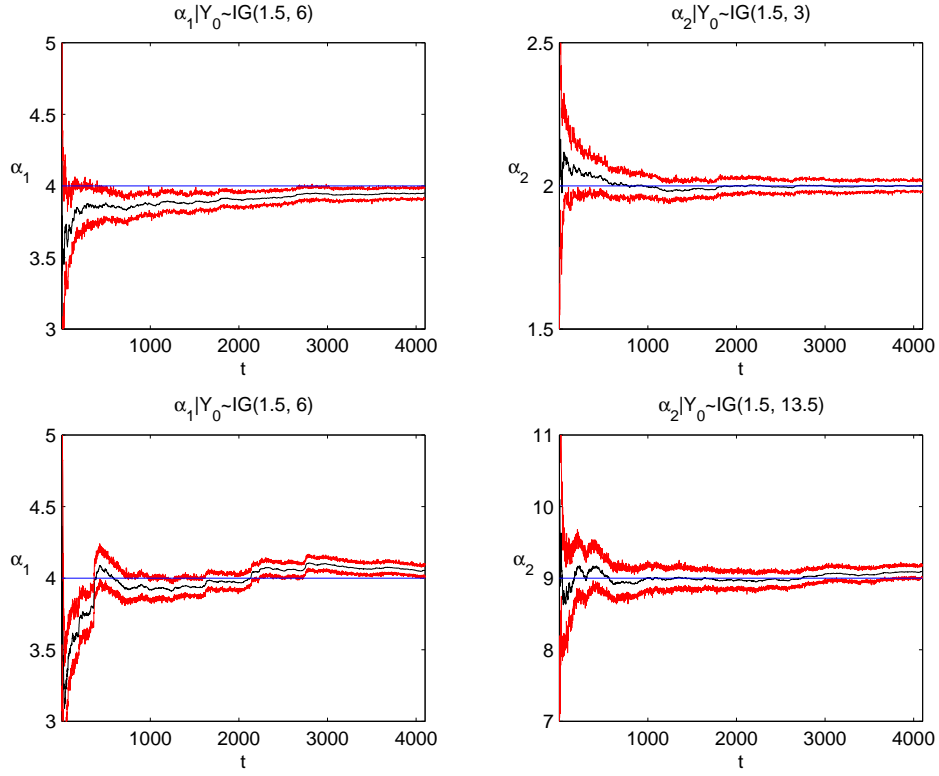


Figure 6: Estimates of temporally and spatially varying scale factors. Each panel shows estimates of the scale factor along with the corresponding 95% empirical credible intervals for the first  $T = 4000$  assimilation cycles. The top panels show the results for temporally varying scale factors, and the bottom panels depict results for spatially varying factors ( $\hat{\alpha}_1$ -left;  $\hat{\alpha}_2$ -right).

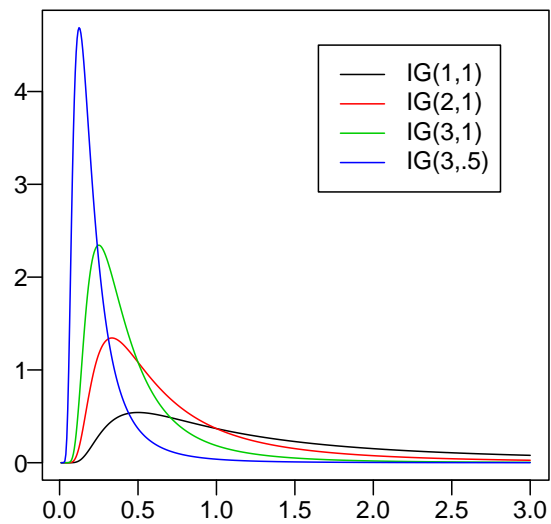


Figure 7: Inverse gamma densities for various parameter values.